






PISCOLA: a data-driven transient light-curve fitter

Tomás E. Müller-Bravo¹^{*}, Mark Sullivan¹, Mathew Smith^{1,2},
Chris Frohmaier¹, Claudia P. Gutiérrez^{1,3,4}, Philip Wiseman¹, Zoe Zontou¹

¹*School of Physics and Astronomy, University of Southampton, Southampton, Hampshire, SO17 1BJ, UK*


²*Univ Lyon, Univ Claude Bernard Lyon 1, CNRS, IP2I Lyon / IN2P3, IMR 5822, F-69622, Villeurbanne, France*

³*Finnish Centre for Astronomy with ESO (FINCA), FI-20014 University of Turku, Finland*

⁴*Tuorla Observatory, Department of Physics and Astronomy, FI-20014 University of Turku, Finland*

Accepted XXX. Received YYY; in original form ZZZ

ABSTRACT

Forthcoming time-domain surveys, such as the Rubin Observatory Legacy Survey of Space and Time, will vastly increase samples of supernovae (SNe) and other optical transients, requiring new data-driven techniques to analyse their photometric light curves. Here, we present the ‘Python for Intelligent Supernova-COSmology Light-curve Analysis’ (PISCOLA), an open source data-driven light-curve fitter using Gaussian Processes that can estimate rest-frame light curves of transients without the need for an underlying light-curve template. We test PISCOLA on large-scale simulations of type Ia SNe (SNe Ia) to validate its performance, and show it successfully retrieves rest-frame peak magnitudes for average survey cadences of up to 7 days. We also compare to the existing SN Ia light-curve fitter SALT2 on real data, and find only small (but significant) disagreements for different light-curve parameters. As a proof-of-concept of an application of PISCOLA, we decomposed and analysed the PISCOLA rest-frame light-curves of SNe Ia from the Pantheon SN Ia sample with Non-Negative Matrix Factorization. Our new parametrization provides a similar performance to existing light-curve fitters such as SALT2. We further derived a SN Ia colour law from PISCOLA fits over ~ 3500 to 7000\AA , and find agreement with the SALT2 colour law and with reddening laws with total-to-selective extinction ratio $R_V \lesssim 3.1$. 

Key words: software: data analysis – supernovae: general – cosmology: observations – distance scale

1 INTRODUCTION

The next decade will see a rapid increase in the number and quality of extragalactic optical transients observed in wide-field, time-domain sky surveys. Detected objects will include classical supernova types, as well as new and exotic transients evolving on a wide-range of timescales. The scientific analyses of these transients will almost always require the fitting of the observed transient light-curves, followed by an estimation of their rest-frame properties, luminosities and timescales. Further, the study of the demographics of the populations these events are drawn from will require the automated fitting of many thousands of transients.

In this paper, we present PISCOLA, a data-driven transient light-curve fitter. Other than practical considerations such as ease of use, there are three main advantages of using PISCOLA for light-curve fitting. The first is that PISCOLA uses Gaussian processes for the light-curve fitting, and thus is not limited by an underlying light-curve model or tem-

plate, and returns the intrinsic light-curve of a transient rather than a template fit. Secondly, other than an spectral energy distribution (SED) that is used for K-corrections, PISCOLA is agnostic to the type of transient being fit. Thus, it can in principle be used to estimate the rest-frame light-curves of any optical transient. Finally, as there is no underlying template, the intrinsic rest-frame light-curves/colours of the transient being fit can be measured across all phases of the transient’s evolution: PISCOLA has no assumed colour evolution or colour law.

As a proof-of-concept of PISCOLA, in this paper we demonstrate the application of PISCOLA to type Ia supernovae (SNe Ia). SNe Ia are the result of a runaway thermonuclear explosion of carbon-oxygen white dwarf star (Hoyle & Fowler 1960; Woosley et al. 1986). The resulting light curves are powered by the radioactive decay of ^{56}Ni (e.g., Colgate & McKee 1969), reaching a peak luminosity about 20 days after explosion (e.g., Riess et al. 1999b). Although the exact triggering mechanism and configuration of their progenitor systems remains controversial (see review of Jha et al. 2019), empirically SNe Ia appear simi-

* E-mail: t.e.muller-bravo@soton.ac.uk

lar in brightness with an intrinsic root-mean-square (r.m.s.) scatter in their peak *B*-band magnitudes of around 0.3 mag (e.g., Cadonau et al. 1985; Tammann & Leibundgut 1990). Correlations between their peak luminosity, the decline rate of their light curves and their optical colours allow them to be further standardised (e.g., Pskovskii 1977; Phillips 1993; Tripp 1998; Phillips et al. 1999; Kattner et al. 2012), reducing their observed r.m.s. scatter to around ~ 0.15 mag or 7–8 per cent in distance (e.g., Betoule et al. 2014). SNe Ia have therefore assumed a key role as cosmological distance estimators (e.g., Kowal 1968; Riess et al. 1998; Perlmutter et al. 1999; Riess et al. 2011; Betoule et al. 2014).

Cosmological studies of SN Ia have always used some form of light-curve fitting (e.g., Leibundgut et al. 1991; Riess et al. 1995, 1996; Perlmutter et al. 1997; Jha et al. 2007; Guy et al. 2007; Conley et al. 2008; Burns et al. 2011), in order to interpolate their peak magnitudes and perform a K-correction from observed magnitudes to the rest-frame. These light-curve fitters traditionally consist of a (multi-colour) time-series light-curve template or SED built from a sample of well-observed, and often low-redshift, SNe Ia. However, although these fitters have worked well for cosmological applications, they have some limitations that a data-driven approach may not.

Firstly, current template-based light-curve fitters are limited by the choice of light-curve parameters to be extracted and do not have the flexibility to easily extract further information from the SNe (for example, colour evolution) to understand their physics. Secondly, template-based fitters typically only work well where they are trained: typically on samples of ‘normal’ SNe Ia with optical data, and excluding (for example) SN1991T- and SN1991bg-like events (Filippenko et al. 1992b,a; Ruiz-Lapuente et al. 1993) or wider wavelength ranges such as the near-infrared (IR). Finally, alternative standardisations with different light-curve parameters can be explored: uncertainties due to the parametrizations used are currently an important factor in the systematic uncertainty budget of SN Ia cosmological analyses (e.g., Abbott et al. 2019).

This paper is organised as follows. In Section 2 we introduce and describe our new light-curve fitter. In Section 3 and 4 we test our fitting code on samples of SNe Ia, both simulated and real. In Section 5 we then present an application of PISCOLA for SN Ia distance estimation, and explore alternative standardisations that PISCOLA can provide. A simple analysis of the SN Ia colour law derived from our light-curve fitter is shown in Section 6, and we summarise in Section 7. Throughout, we assume a flat Λ CDM cosmology with $H_0 = 70 \text{ km s}^{-1} \text{ Mpc}^{-1}$, and $\Omega_M = 0.3$.

2 PISCOLA

In this section, we introduce the light-curve fitter PISCOLA (Python for Intelligent Supernova-COSmology Light-curve Analysis). The code is open source¹ and written in PYTHON 3, a language widely used in the astronomy community. PISCOLA is intended to be simple and transparent in operation, allowing the user the opportunity to understand

every step of the light-curve fitting and correction process. The full documentation can be found online².

We introduce PISCOLA in the context of fitting SN Ia light curves and subsequently in SN Ia distance estimation, but stress that PISCOLA is flexible and can be used to fit the light curves of any optical transient. In Fig. 1, we show the main steps in PISCOLA, described in the following sections.

2.1 The light-curve model

PISCOLA uses Gaussian processes (GPs), a Bayesian method, to model the SN light curves by assuming that the data distributes as a multi-variate Gaussian distribution, which is the case for most astronomical observations. Unlike many other SN light curve fitters, there is no underlying SN template or model that is directly fit to the observed light curves; the observer-frame modelling is entirely data driven. This implies that our method needs better quality data than template-driven methods to reliably fit SNe Ia: it is not suited to poorly sampled data. In Section 3, we determine the quality of the data necessary for PISCOLA to produce reliable results.

GP is an excellent tool for data regression, providing a more robust regression than polynomials as they can be undesirably global (splines are just a special case of GP regression). GPs have also been previously used in the context of SN and SN-like light curves, with excellent results (e.g., Kim et al. 2013; de Jaeger et al. 2017; Inseerra et al. 2018; Angus et al. 2019; Pursiainen et al. 2020; Wiseman et al. 2020).

A GP model is defined by a mean and a covariance function/kernel. There are different types of kernels (e.g., ‘squared exponential’, ‘rational quadratic’, ‘periodic’, the ‘Matérn family’, etc.), each with its own set of hyperparameters (see Appendix A for more details). PISCOLA makes use of GEORGE (Ambikasaran et al. 2016), a GP implementation in PYTHON that allows the user to choose between several well-known kernels, together with the SCIPY and LM-FIT packages for the optimisation routines. Our code implements three kernels: squared exponential, Matérn-3/2 and Matérn-5/2. The Matérn family is broadly used in different areas of research as it is effective at describing different physical processes (Rasmussen & Williams 2006), such as light curves, while the squared exponential kernel provides a smoother fit. We use different kernels at different points in our light-curve fitting process.

The fitting of the SN multi-colour observer-frame light curves is performed in two dimensions, as a function of time and wavelength. All three hyperparameters are optimised at this stage: length scales for the time and wavelength axes, and the variance (see Appendix A for details about these hyperparameters). We fit in magnitude space as the logarithmic scale provides smoother and more accurate GP fits than using fluxes directly; when fitting in flux space, any large differences in flux between different filters can produce fits with under-estimated peaks (this happens for different choices of kernel).

There are two disadvantages of fitting in magnitude space. The first is that non-detections are not used as they

¹ <https://github.com/temuller/piscola>

² <https://piscola.readthedocs.io/en/latest/>

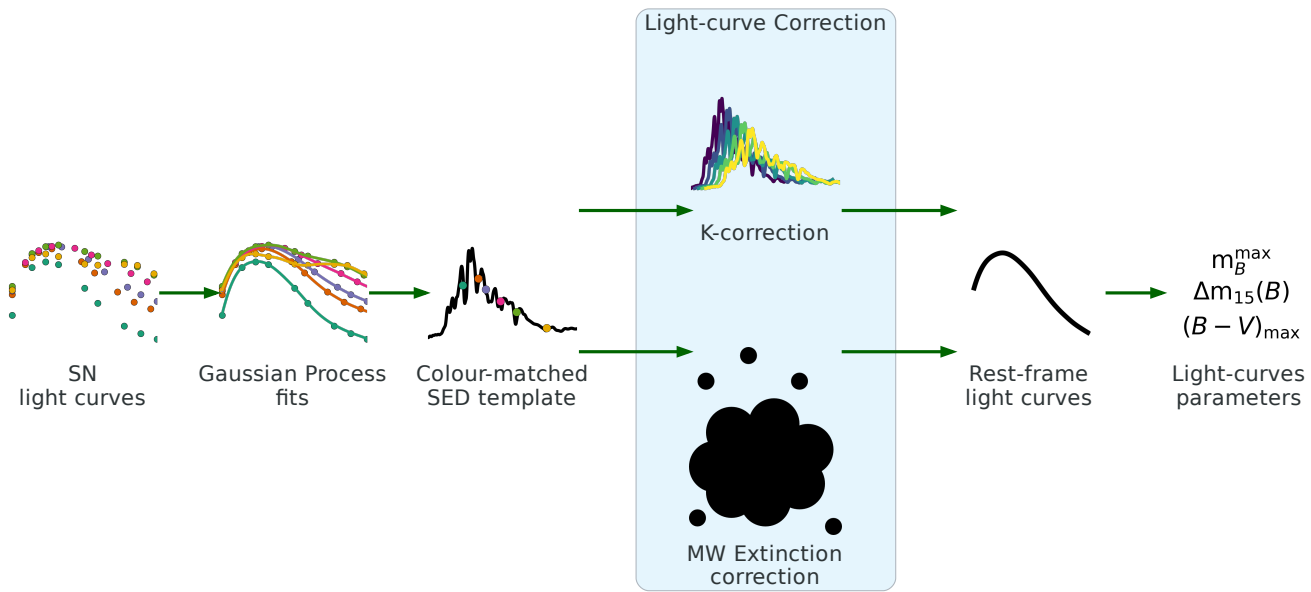


Figure 1. Flowchart of the main steps in the PISCOLA process. Gaussian processes are used to fit the SN light-curves as a function of wavelength and time (Section 2.1). An SED is then ‘warped’ to match the observed SN colours, corrected for redshift (K -correction) and then Milky Way dust extinction (Section 2.2). Finally, the required rest-frame light curves are obtained and, optionally, light-curve parameters estimated from those light curves.

cannot be represented correctly in magnitude space. This also means that data around peak luminosity will have more weight than data in the tails of the light curves, for measurements with the same uncertainty in flux space, during the regression. We note that it is the data around peak that interests us most in the light-curve fitting for the applications in this paper. A second disadvantage is that lower S/N observations have asymmetric errors, which are an issue for GP regression. However, these can be removed by masking out these observations (see Section 3), although this reduces the amount of useful data.

The two-dimensional regression provides more robust results than fitting in one dimension at a time, i.e., flux as a function of time independently for each band, as it uses information from multiple bands to cover time gaps in different filters allowing an improved interpolation (see Appendix A1). An example of the fitting applied to a high-redshift SN Ia *griz* light curve is shown in Fig. 2.

A Matérn-5/2 kernel is the default option in PISCOLA as it produces flexible, yet smooth fits. In Fig. 2, the fit to the redder z -band is good, despite having larger uncertainties. In addition, although only the i -band has relatively good coverage of the tails of the light curve both pre- and post-peak, the two-dimensional GP model allows an informative extrapolation of the time axis of the grz -band light curves. For comparison, an example with a one-dimensional GP model is shown in Fig. A4 and described in Appendix A1.

A further advantage of fitting in two dimensions is that more accurate estimations of the time of peak luminosity in a given rest-frame filter are obtained compared to 1D fits in individual observer-frame filters. The time of peak luminosity in a given band is estimated from the GP fit by calculating the time at which the derivative becomes zero at the effective wavelength of the desired band. For the case of

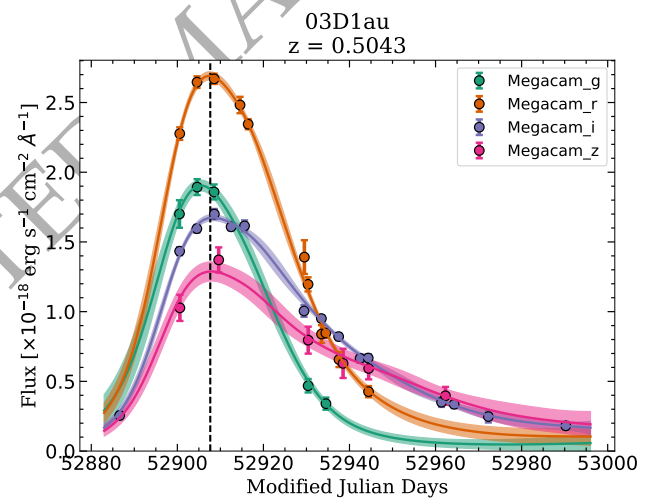


Figure 2. Two dimensional GP model of the observer-frame multi-colour light curve of the SN Ia SNLS-03D1au using a Matérn-5/2 kernel. The original data, from the Supernova Legacy Survey (Astier et al. 2006), were observed in the Canada–France–Hawaii Telescope MegaCam *griz* filters and are shown as circles with uncertainties. The solid lines show the mean of the GP model, and the shaded areas represent one standard deviation (1σ). The vertical dashed-black line marks the estimation of the time of rest-frame B -band maximum light.

the rest-frame B band, we denote the time of peak luminosity as parameter t^{\max} . This is important in many applications, but particularly in the light-curve correction processes for SNe Ia in cosmology.

Moreover, PISCOLA can produce excellent results with relatively well-sampled SNe Ia, such as those at low red-

shift ($z \lesssim 0.1$), as it is able to fit bands outside the useful ranges of other light-curve fitters. An example is shown in Appendix B. This is of great importance in many areas, for example in the estimation of the Hubble constant, measurement of peculiar velocities in the local universe, and the extraction of astrophysics from the SNe Ia light curves, to name just a few applications.

2.2 The spectral energy distribution model

Although no light-curve template is used in PISCOLA, a (time-dependent) SED is still required to K-correct the observer-frame light curves to rest-frame filters. Many such templates are available for different SN types (e.g., Nugent et al. 2002; Blondin & Tonry 2007; Hsiao et al. 2007; Rodney & Tonry 2009; Liu & Modjaz 2014; Vincenzi et al. 2019). We colour-adjust (or ‘mangle’) the SED as a smooth function of wavelength so that it reproduces the colours of the observed light curves at each epoch, as estimated by the GP model (see also Hsiao et al. 2007; Conley et al. 2008). The result is a colour-matched SED, i.e., a SN time-series SED which reproduces the observed light curves of the SN being fit.

The colour-matching is performed by multiplying the SED by a wavelength-dependent function at the desired phases (with respect to t^{\max}), after which the SED reproduces the observed light-curve when integrated through the observed filters. This wavelength-dependent function is often represented by a spline in the literature, but PISCOLA makes use of GPs with (by default) a squared exponential kernel, resulting in a smooth function and a natural way of propagating uncertainties in the process. The steps for developing the colour-matching function for each SN are:

- (i) An observer-frame SED is constructed from the target’s redshift;
- (ii) Milky Way (MW) dust extinction is applied to the SED;
- (iii) The observed filters are used to calculate model fluxes from the SED;
- (iv) The ratios between the actual fluxes of the observed SN and the model fluxes from the SED are calculated together with their respective effective wavelengths;
- (v) A wavelength-dependent multiplicative function is calculated by modelling the flux ratios with GPs using a squared exponential kernel.

The SED template is multiplied by the multiplicative function to produce a colour-matched SED. An example of a typical multiplicative function and SED template is shown in Fig. 3.

As a final step, the colour-matched SED is corrected for MW extinction and shifted to the rest-frame to obtain a final SED model. PISCOLA corrects for MW dust extinction using the Schlafly & Finkbeiner (2011) dust maps and a Fitzpatrick (1999) extinction law as default. Other implementations of the dust maps (Schlegel et al. 1998) and extinction laws are also available. The SFDMAP and EXTINCTION PYTHON packages are used to perform the corrections. PISCOLA does not estimate, or correct for, any other sources of extinction. This procedure is repeated for as many phases as desired for each SN, depending on the data coverage.

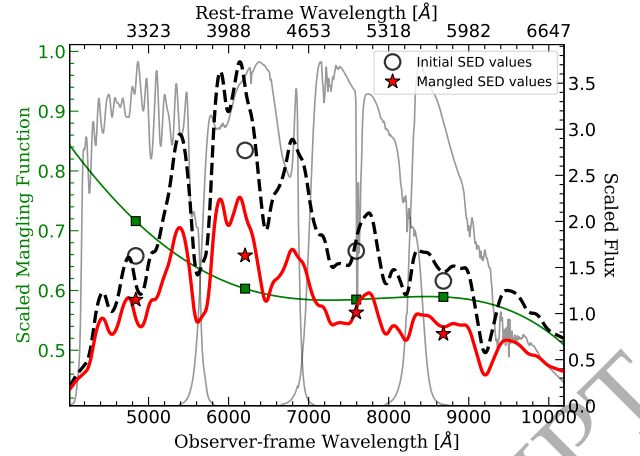


Figure 3. Mangling function (solid green line) for SNLS-03D1au at t^{\max} . The green squares represent the flux ratios of the different bands centered at their respective effective wavelengths. The initial SED (dashed black line) is compared to the SED multiplied by the mangling function (solid red line). The initial (black open circles) and mangled (red stars) SED fluxes are also shown. The transmission functions of the filters used are plotted in grey. The scaling is arbitrary for visualisation.

The product of this process is a rest-frame time-series SN SED that reproduces the observer-frame light curves of the SN being modelled. From this, rest-frame light curves in any required band can be calculated within the wavelength and phase limitations of the input data (together with any desired light-curve parameters based on those rest-frame light curves). This process is agnostic to the physical type of transient being fit, under the assumption that a physically relevant time series SED is used in the colour-matching process.

We note one subtlety: the procedure described above depends on the initial estimate of t^{\max} . This initial estimate can be improved using the rest-frame light curves. If the initial and improved estimates of t^{\max} are not consistent, the light-curve fitting process can be repeated with the new estimate until a satisfactory convergence is reached.

In the remainder of this section, and to assess the performance of PISCOLA, we discuss the specific use of PISCOLA to the fitting of SN Ia light curves. Here we focus on traditional SN Ia rest-frame light-curve parameters: the rest-frame B -band peak apparent magnitude (m_B^{\max}), the decline in magnitudes in the 15 days following t^{\max} ($\Delta m_{15}(B)$), and the $B-V$ colour at t^{\max} ($(B-V)_{\max}$). In Fig. 4, we show an example of a resulting rest-frame B -band light curve of a SN Ia.

2.3 Application to SNe Ia

Several aspects of the light-curve fitting process can be affected by choices in the analysis, including the MW extinction law, the SED time-series template and the exact rest-frame filters. In Section 4, when we will compare PISCOLA with results from the SALT2 SN Ia light-curve fitter to establish its performance, an accurate comparison will require these choices to be consistent.

SALT2 is an empirical SN Ia model, used in several ma-

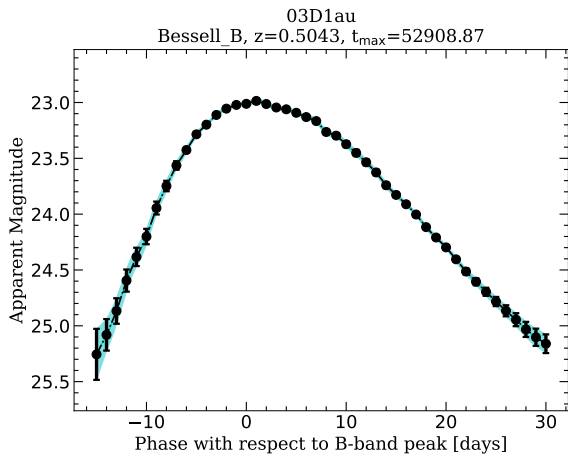


Figure 4. Rest-frame B -band light curve of SNLS-03D1au. In this case, observations with Megacam-*griz* filters were used to construct the B -band light curve.

for cosmological analyses (Betoule et al. 2014; Scolnic et al. 2018b; Abbott et al. 2019), trained on a spectrophotometric sample of SNe Ia. It describes the flux of a SN Ia with the following functional form,

$$F(SN, p, \lambda) = x_0 \times [M_0(p, \lambda) + x_1 M_1(p, \lambda) + \dots] \times \exp[c \times CL(\lambda)], \quad (1)$$

where p is the rest-frame time with respect to B -band peak, λ is the rest-frame wavelength, M_0 is the average spectral sequence, M_1 is an additional component that describes further variability, and CL represents the average colour variation law of a SN Ia (Section 6). The contribution of higher-order components is less significant and not used. The terms x_0 , x_1 and c are the SALT2 light-curve parameters, where $-2.5 \log_{10}(x_0) + 10.635 = m_B^{\max}$, x_1 is a measurement of the stretch of the light curve, and c is the colour of the SN.

We ensure our choices follow Scolnic et al. (2018b) as far as possible:

- We adopt the extinction law from Fitzpatrick (1999);
- The SN Ia SED time-series template is the M_0 component from the SALT2 model (`salt2_template_0.dat`). We do not include the SALT2 x_1 -dependent component in the SED as PISCOLA does not measure this parameter (i.e., a template with $x_1 = c = 0.0$ is used);
- We use the same filter transmission functions for the observed light curves as Scolnic et al. (2018b);
- We use the same photometric calibration system for magnitude systems from different surveys (e.g., AB using BD +17°4708 (Bohlin & Gilliland 2004) as primary standard);
- We use the Bessell filters (Bessell 1990), shifted to match Landolt (1992) observations, as included in SALT2 (see Betoule et al. 2014, and Appendix A of Conley et al. 2011), to estimate the rest-frame light-curve parameters.

In principle, the use of the M_0 SED but neglecting the spectral differences introduced by x_1 , will lead to small differences in the K -corrections when compared to SALT2. For example, when including the x_1 -dependent component (assuming the x_1 values from Scolnic et al. 2018b), the differ-

ences in the rest-frame B -band light curve are on average $\lesssim 0.01$ mag around peak and $\lesssim 0.03$ mag at earlier or later phases. This sets a natural limit to the accuracy of these tests.

3 TESTS WITH SIMULATIONS OF SNE IA

We test PISCOLA with extensive simulations using the SuperNova ANALysis software (SNANA; Kessler et al. 2009b), version v10_75c. We focus our tests on the effects of cadence, the mean time between consecutive observations of an object in the same filter, and observational uncertainties or, equivalently, signal-to-noise ratio (S/N).

SNANA simulations simulate transient surveys and produce light curves of simulated transient events, accounting for the survey observing pattern, limiting magnitudes, pointing on the sky, and so forth. For testing PISCOLA, we simulate Pantheon-like SN Ia samples (Scolnic et al. 2018b), representing one of the most comprehensive compilations of SNe Ia. Pantheon comprises many different SN Ia surveys: the Harvard-Smithsonian Center for Astrophysics (CfA) surveys 1–4 (Riess et al. 1999a; Jha et al. 2006; Hicken et al. 2009a,b, 2012), the Carnegie Supernova Project (CSP; Contreras et al. 2010; Folatelli et al. 2010; Stritzinger et al. 2011), the Sloan Digital Sky Survey SN Survey (SDSS; Frieman et al. 2008; Kessler et al. 2009a; Sollerman et al. 2009; Sako et al. 2018), the Supernova Legacy Survey (SNLS; Astier et al. 2006; Guy et al. 2010), the Panoramic Survey Telescope and Rapid Response System 1 (Pan-STARRS1 or PS1) Medium Deep Survey (Rest et al. 2014; Scolnic et al. 2014) and different *Hubble Space Telescope* (HST) surveys, including: the Supernova Cosmology Project (SCP; Suzuki et al. 2012), the Great Observatories Origins Deep Survey (GOODS; Riess et al. 2007), and the Cosmic Assembly Near-infrared Deep Extragalactic Legacy Survey + Cluster Lensing And Supernova Survey with Hubble (CANDELS+CLASH; Graur et al. 2014; Rodney et al. 2014; Riess et al. 2018).

We denote the compilation of the CfA1–4 and CSP surveys the ‘low- z ’ sample (or survey) as they only contain SNe Ia at $z < 0.1$. For a more detailed description of the surveys, see Scolnic et al. (2018b) or their respective references. We use the SALT2 light-curve parameters for these SNe Ia from Scolnic et al. (2018b)³.

All simulations have the same main characteristics. We restrict our simulations to the ground-based samples (low- z , SDSS, SNLS and PS1) as PISCOLA does not perform well with the less well-sampled light curves that are typical from the *HST* surveys. Simulated SNe Ia are based on the SALT2 model from Betoule et al. (2014, version 2.4), also used by Scolnic et al. (2018b). The SNe are simulated by selecting random x_1 and c values drawn from the distributions of Scolnic & Kessler (2016), which are survey dependent based on several factors such as Malmquist bias; for more details, see Scolnic & Kessler (2016). Approximately 500 SNe are simulated for each SN sample (low- z , SDSS, SNLS and PS1) and for each individual test. In other words, each simulated

³ <https://archive.stsci.edu/prepds/ps1cosmo/>

Pantheon-like sample contains a total of ~ 2000 SNe Ia. We discuss this choice in Section 3.4.

The comparison between the fits and the simulations are based on m_B^{\max} , i.e., we focus on the residual between the simulated m_B^{\max} value and the estimated value obtained with PISCOLA. In this paper, we will refer to this as $\Delta m_B^{\max} \equiv m_{B, \text{PISCOLA}}^{\max} - m_{B, \text{simulations}}^{\max}$. We fit the SNe in an automated way with the default GP kernels and over -15 to $+30$ days with respect to t^{\max} (default phases). In addition, we require every SN to have a sufficient coverage of the peak: at least one data point in any band over -7 and 0 days, one from 0 to $+7$ days, and one from -3.5 and $+3.5$ days (this final constraint may overlap with one or both previous constraints). We also mask out observations with a $S/N \leq 5$ to prevent poor-quality fits and asymmetric errors. We note that applying cuts in S/N does not bias the estimation of m_B^{\max} or the observed peaks in the fitted light curves (this was tested on real data and simulation from the following sections), although it could possibly bias the fits at earlier or later epochs. However, for the analysis in this work, we are mainly concerned about the data around peak. Note that the cuts applied are different, but more stringent than those applied in other analyses using SALT2 (e.g., Scolnic et al. 2018b), discarding more SNe.

These constraints reduce the sample of SNe, but they ensure a relatively good quality light curve data required by GP interpolation (to be tested in the following sections). The chosen values are somewhat arbitrary, but small changes do not significantly alter the results. We note that, although template-driven fitters can fit SNe with different S/N , they may introduce biases as, in the case of low S/N , the fitters will mainly retrieve the light-curve parameters of the templates being used.

3.1 SN Ia Pantheon sample simulations

The first step is to test PISCOLA with a Pantheon-like simulation. The comparison between the PISCOLA-measured and simulated m_B^{\max} values is shown in Fig. 5.

We discard some SNe at different stages of the fitting process for several reasons:

- (i) The observer-frame wavelength coverage does not cover the rest-frame B -band at the redshift of the SN (e.g., SNe at high redshift);
- (ii) The temporal coverage does not allow an initial estimation of t^{\max} (e.g., many low- z do not have data prior to the B -band peak);
- (iii) PISCOLA is unable to estimate an accurate B -band light curve. This happens when PISCOLA does not converge to an estimation of t^{\max} ; e.g., several low- z SNe have their B -band peak only partially covered, and t^{\max} can fall close to the limits of the coverage, producing a failure to estimate a new peak;
- (iv) The light-curve does not satisfy the additional constraints on peak coverage (Section 3), mainly due to a combination of low cadence and/or masking of low S/N data;
- (v) Visual inspection reveals poor PISCOLA fits, mainly caused by remaining poor data quality.

The first three reasons are automated by PISCOLA. Discarding SNe after visual inspection is only feasible due to

Table 1. Weighted mean, uncertainty on the weighted mean and weighted standard deviation of Δm_B^{\max} for a Pantheon-like simulation (Sec. 3.1).

survey	weighted mean (mmag)	error on the weighted mean (mmag)	weighted standard deviation (mmag)
low- z	1	1	24
SDSS	4	2	34
SNLS	3	2	29
PS1	7	2	35

the relatively low number of SNe. This is a limitation, and for future implementations we aim to automate it with a statistically-motivated metric, which is the aim of this set of tests. We note that future surveys (e.g., LSST), will produce higher quality observations with lower uncertainties and better cadences.

We successfully fitted (obtained m_B^{\max}) ~ 50 per cent of the SNe. This is a low number compared to other light-curve fitters, but highlights the relatively poor data quality of some historical surveys that SNANA simulates. Stages (ii) and (iv) are the main causes of discarding SNe (see Section 4). We note that these discarding reasons do not strongly depend on redshift but mainly on the observing strategy of each survey.

In Table 1, we show the values of Δm_B^{\max} . There are no significant (i.e., $< 3\sigma$) deviations from $\Delta m_B^{\max} = 0.0$ mag for all surveys, except for PS1 at 3.5σ significance but small (< 0.01 mag) deviation. This is mainly caused by a few SNe with underestimated m_B^{\max} and small uncertainty. Additionally, Fig. 5 shows that some low- z SNe have underestimated m_B^{\max} values, while some SNLS SNe have the opposite. This is most likely due to a combination of cadence and S/N and is further investigated in the following section. Generally, we conclude that PISCOLA is successful at retrieving accurate m_B^{\max} values at < 0.01 mag.

3.2 Effect of observational cadence

As PISCOLA is a data-driven fitting method, the cadence is important as the GP model has no prior for the true shape of the SN light curves. We therefore simulated a set of Pantheon-like samples with cadences between 1 and 10 days in steps of 1 day, and took the 1-day cadence simulation as our benchmark. Observations are equally spaced for all bands (e.g., simultaneous/same-day *ugriz*-bands observations every x days) and with random characteristics drawn from survey-dependent distributions. We then estimate the reliability of PISCOLA as the cadence becomes poorer. The results are shown in Fig. 6.

We successfully fitted (obtained m_B^{\max}) ~ 75 – 85 per cent of the SNe in each of the 1- to 7-day cadence simulations, and ~ 60 – 75 per cent for the 8- to 10-day cadence simulations (SALT2 fits $\gtrsim 98$ per cent of the simulated SNe in all cases). The unsuccessful SNe are mainly due to low- S/N light curves in the SDSS, SNLS and PS1 simulated samples and B -band peak not well covered (e.g., data only after peak) in the case of low- z objects. SNe were discarded for the same reasons as in Section 3.1. At relatively high cadence ($\lesssim 7$ d) PISCOLA is accurate when estimating m_B^{\max} for the different surveys




Figure 5. $\Delta m_B^{\max} \equiv m_B^{\max, \text{PISCOLA}} - m_B^{\max, \text{simulations}}$ for a Pantheon-like simulation. Each point represents a single simulated SN, colour-coded according to the legend. The uncertainties are 1σ and are taken from the PISCOLA fits. The golden hexagons represent the weighted mean in bins of 1 mag with their respective uncertainties (1σ).

(i.e., $\gtrsim 3\sigma$ deviations from $\Delta m_B^{\max} = 0.0$ mag; see bottom panel of Fig. 6). However, the accuracy is improved for SNLS and PS1 compared to low- z and SDSS samples, given their superior S/N and/or rest-frame cadence.

As redshift increases, the rest-frame cadence increases for fixed observer-frame cadence. This gives improved light-curve coverage, but with a trade-off of S/N. We examined Δm_B^{\max} versus redshift, but see no clear trend. The surveys at different redshifts also generally perform similarly as a function of cadence. We find that PISCOLA accurately retrieves m_B^{\max} ($\lesssim 3\sigma$ deviations from $\Delta m_B^{\max} = 0.0$ mag) for cadences similar to those of most high-redshift transient surveys, which have a typical cadence of about 7 days in the observer frame (e.g., Dark Energy Survey).

However when the cadence is $\gtrsim 7$ days, the performance of PISCOLA on the low- z simulated sample is quantitatively different to the performance on the higher-redshift simulations. At low- z , PISCOLA under-predicts m_B^{\max} (Δm_B^{\max} is positive). This is because the peaks of the light curves are smoothed out: the GP model has no information about their true shape and thus does not recover the peak given a lack of information. In the other simulations (SDSS, SNLS and PS1), the opposite trend is seen: Δm_B^{\max} is systematically negative at low cadences. This is due to the lower S/N in these data, despite the higher rest-frame cadence.

SNe with the lowest S/N, which would normally produce slightly under-predicted m_B^{\max} values (see Fig. A5), are not successfully fit, and therefore, not included in the comparison, producing the effect of an apparent over-prediction

of m_B^{\max} for these samples instead of just an increase in scatter. The masking of low-S/N data in part produces a similar effect. This explains some of the discrepancies seen for some SNLS SNe in Section 3.1. SNe with high S/N, but relatively low cadences, can have their m_B^{\max} values slightly over-predicted due to measurements with low uncertainty at each side of the light-curve peak, which can produce sharper peaks (top left panel of Fig. A5). However, this is not the general case (see the top left panel of Fig. 6) as it very much depends on the time of the observations with respect to the actual light-curve peak and the apparent luminosity of the SN.

These effects can be reduced by tightening the constraints of the observations around peak, although this would in turn reduce the number of usable SNe. In practice, real observations are not evenly spaced, so these tests just provide a more general idea of how well PISCOLA performs.

3.3 Effect of observational uncertainties

The GP interpolation depends on the uncertainties of the observations. Thus, we tested the effect of observational uncertainties by simulating two Pantheon-like samples with half and twice the observational uncertainties (σ_{obs}) of the data in the original Pantheon sample, with a fixed cadence of 7 days.

The weighted mean and weighted standard deviation of Δm_B^{\max} for the 7-day cadence simulations with original, half

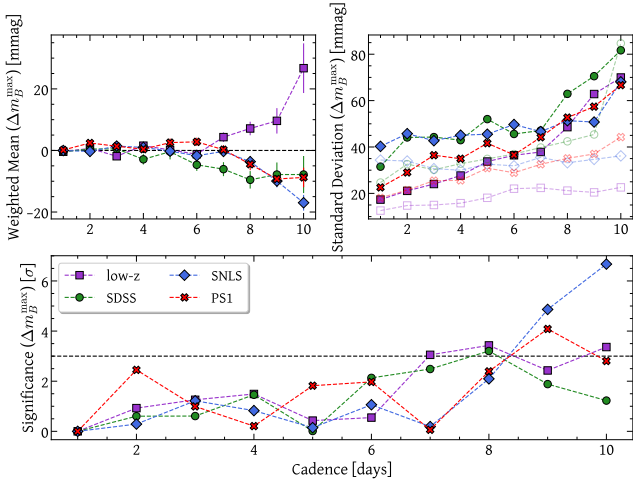


Figure 6. Weighted mean with uncertainty on the weighted mean (**top left** panel) and standard deviation (**top right** panel) of Δm_B^{\max} for simulations with cadences between 1 and 10 days. The standard deviation of Δm_B^{\max} for SALT2-measured m_B^{\max} values, i.e., m_B^{\max} , SALT2 - m_B^{\max} , simulations, are also shown for comparison (transparent symbols in the **top right** panel). Measured significant deviations of $\Delta m_B^{\max} = 0.0$ mag in units of σ are also shown (**bottom** panel). The horizontal line in the bottom panel marks a significance of 3σ .

and twice the observational uncertainties are shown in Table 2. From this comparison, we can see that a deviation from $\Delta m_B^{\max} = 0.0$ mag (at $\sim 3\sigma$) for the SDSS sample is persistent for different values of σ_{obs} , which may be caused by a generally low S/N despite the $S/N \geq 5$ selection used. For the low- z sample, a significant (3.5σ) deviation from $\Delta m_B^{\max} = 0.0$ mag (PISCOLA under-predicts m_B^{\max}) is only observed for half σ_{obs} given the relatively low cadence, despite the high S/N, as discussed at the end of Section 3.2. In the case of SNLS, a significant deviation (4.5σ) in the estimation of m_B^{\max} is only seen for twice σ_{obs} , while for PS1 no significant deviations are observed.

Despite having some disagreement in the estimation of m_B^{\max} , these are all < 0.01 mag: PISCOLA is in good agreement with established light-curve fitters. However, the characteristics of surveys like SNLS and, especially, PS1, are ideal for fitting SNe with PISCOLA as they have relatively high S/N and the high- z observations allow a good light-curve coverage due to the relatively high rest-frame cadence. Surveys with the characteristics of the low- z sample require high cadences ($\lesssim 6$ days). Future surveys, such as the LSST, will produce high-quality data with good S/N and cadence (e.g., Lochner et al. 2018; Scolnic et al. 2018a), overcoming some of the limitations found in this work, thus allowing PISCOLA to produce reliable fits.

These comparisons help validate our code. However, we note the caveat that the simulations are based on the SALT2 model, which differs from other light-curve fitters. Nonetheless, the results of the tests performed throughout this section are promising as they help establish the reliability of PISCOLA. The level of discrepancies uncovered in this section would only be important in the use of SNe Ia as cosmological probes; in other astrophysical applications the level of disagreement is not significant.

Table 2. Weighted mean, uncertainty on the weighted mean and weighted standard deviation of Δm_B^{\max} for the 7-day cadence simulations with initial, half and twice the observational uncertainties (σ_{obs} ; see Sec. 3).

survey	weighted mean (mmag)	error on the weighted mean (mmag)	weighted standard deviation (mmag)
Initial simulation			
low- z	4	2	26
SDSS	-6	2	39
SNLS	0	2	29
PS1	0	1	30
half σ_{obs} simulation			
low- z	7	2	26
SDSS	-8	2	34
SNLS	4	2	29
PS1	1	1	29
twice σ_{obs} simulation			
low- z	4	2	28
SDSS	-8	3	45
SNLS	-9	2	36
PS1	-4	2	36

3.4 Computational considerations

PISCOLA is a computationally intensive light-curve fitter. Fitting the light curves is fast (of the order of seconds); however, calculating the mangling function can take longer (of the order of minutes for the default phase between -15 and 30 days with respect to t^{\max}). This is because a fixed GP model is used inside a minimization routine to calculate the mangling function. The length-scale of the GP model is fixed to a value of 20 to produce a smooth function, however, the ratios between the observed and model fluxes (see Section 2.2) are treated as parameters for the minimisation routine. We note that the wavelength axis is divided by 1000 before the minimization routine, and before setting the GP length-scale, to ensure better results by avoiding large numbers, as these are not always properly handled by the routine used. The results are then re-scaled by 1000.

The large covariance between the different bands, given by the Squared Exponential kernel used (see Fig. A1), and the precision required, make the calculation of the mangling function a slow process as the minimisation routine takes longer to converge. Additionally, if the whole light-curve correction process needs to be repeated with an improved estimation of t^{\max} (Section 2.2), the time increases further.

As a result, the number of SNe that can be fit in a ‘reasonable’ amount of time is limited to a few thousands or tens of thousands. For instance, fitting 1000 simulated SNe Ia takes at least ~ 2000 minutes (~ 33 hours) on a single CPU core, depending on various factors. This motivated our choice of ~ 2000 simulated SNe per test, allowing us to fit objects in a reasonable amount of time.

4 COMPARISON WITH SALT2

We next test PISCOLA on real data, using the 1022 SNe Ia from the Pantheon sample (excluding HST objects) for these tests. Unlike the simulations in the previous section, we

Table 3. Weighted mean, uncertainty on the weighted mean and weighted standard deviation of Δm_B^{\max} and $(B - V)_{\max} - c$ for the Pantheon sub-sample used in this work (see Sec. 4).

survey	weighted mean (mmag)	error on the weighted mean (mmag)	weighted standard deviation (mmag)
Δm_B^{\max}			
low-z	10	4	31
SDSS	-26	2	31
SNLS	-8	2	25
PS1	-9	2	34
$(B - V)_{\max} - c$			
low-z	45	1	8
SDSS	35	3	38
SNLS	9	5	27
PS1	4	4	44

have no ‘ground truth’ for these tests, so instead we compare the outputs of PISCOLA against those of SALT2 using the light-curve parameter estimations as given in Scolnic et al. (2018b). We fit the SNe in the same way as we described in the previous section (using default PISCOLA parameters) and applied the same constraints as for the simulations (see Section 3). SNe were discarded at different stages of the fitting process following Section 3.1.

Of the 1022 initial SNe Ia, we obtain successful fits, with m_B^{\max} values, for 620, of which 413 have a $(B - V)_{\max}$ estimate. Table 4 shows a summary of the discarded SNe. The percentage of successfully fitted SNe (~ 60 per cent) is larger than for the Pantheon-like simulation in Section 3.1 (~ 50 per cent), which is explained by the different relative numbers of SNe for the different surveys: in the simulations, all surveys have approximately the same number of SNe.

The outputs from different light-curve fitters are difficult to compare on a SN-by-SN basis due to, for example, different internal calibrations in the fitters producing offsets in some parameters. These are not important for cosmology as long as each code is self-consistent. In this section, we focus on the comparisons of m_B^{\max} and $(B - V)_{\max}$ which are the most directly comparable between SALT2 and PISCOLA.

4.1 B-band peak magnitude comparison

Before comparing m_B^{\max} values, an offset of ≈ 0.27 mag needs to be applied as the version of SALT2 used in Scolnic et al. (2018b) incorporates this global offset. On close inspection, we also found that the MW extinction reddening values adopted by Scolnic et al. (2018b), using the Schlafly & Finkbeiner (2011) dust maps, had some inconsistencies, principally in the low- z sample where most SNe have smaller reddening values than expected. We use the same reddening values as Scolnic et al. (2018b) for comparison purposes.

In Fig. 7, we show the results of the comparison between m_B^{\max} . The weighted average and weighted standard deviation of $\Delta m_B^{\max} \equiv \Delta m_{B, \text{PISCOLA}}^{\max} - \Delta m_{B, \text{SALT2}}^{\max}$ are shown in Table 3. There is a general formal disagreement ($>3\sigma$) for all the surveys except the low- z sample.

The differences are typically small, but there are some

important details. The apparent offset observed for the low- z sample is mainly driven by one SN, SN2004ey (in the [14, 15] mag bin in the left panel of Fig. 7), due to its large discrepancy with SALT2 (m_B^{\max} residual of ~ 0.1 mag), but small uncertainties. If this object is removed, the weighted average of the m_B^{\max} residual for the low- z sample is reduced to 0.005 mag. Using SALT2 to fit this SN, we obtained relatively good fits to the light curves (see Fig. B2), but with a different value of m_B^{\max} compared to that published by Scolnic et al. (2018b), although in better agreement with the PISCOLA value. We do not know the exact source of this difference and do not have any reason from the PISCOLA fit, which results in better residuals than SALT2 (see Fig. B1), to discard this SN.

The discrepancy in m_B^{\max} for the SDSS sample is the largest, expected due to its relatively larger uncertainties. The lower S/N may cause an apparent over-prediction of PISCOLA-measured m_B^{\max} values for these SNe (see discussion in Section 3.2). However, there is also milder disagreement for the SNLS and PS1 surveys, perhaps due to unidentified issues with photometric calibration given the results of Section 3, or due to differences in the SED models between SALT2 and PISCOLA. Despite thoroughly checking the analysis and our tests, we are unable to identify the source of this discrepancy. We note that the differences could be due to unidentified issues in PISCOLA, SALT2 or both.

4.2 Colour comparison

We compare the colour parameters in Fig. 8. Given the data-driven nature of PISCOLA, not every SN with an estimation of m_B^{\max} has sufficient wavelength coverage to also estimate its $(B - V)_{\max}$, particularly at high redshift. We also note that the colour parameters are fundamentally different (see, e.g., Kessler et al. 2013), and thus a detailed comparison is difficult. For example, SALT2 estimates $(B - V)_{\max}$ through the c parameter using information across the 3000–7000 Å range, while PISCOLA makes a direct measurement of $B - V$.

The weighted average and weighted standard deviation for the colour parameter residuals are shown in Table 3. As expected, some differences are seen between the colour parameters, particularly in the low- z and SDSS samples, where PISCOLA measures redder SNe than SALT2. A global offset may be expected, but the discrepancies are different for each survey, possibly implying issues with calibration as for m_B^{\max} , or with the SED model.

In summary for this section, we find some differences between light-curve parameters for SNe Ia measured using PISCOLA and SALT2, some of which were anticipated. For the most straight forward comparison, i.e. m_B^{\max} , the differences between PISCOLA and SALT2 were small. A detailed comparison for colour is more difficult due to slightly varying definitions. Taking the results from this section and from Section 3, we conclude that the performance of PISCOLA is satisfactory and validates our code.

Table 4. Number of supernovae discarded at different stages in the PISCOLA fitting.

Discarding reason	low-z	SDSS	SNLS	PS1	Total	Cumulative number discarded
Poor wavelength coverage ^a	1	3	33	9	46	46
Poor time coverage ^b	58	21	7	11	97	143
Inaccurate results ^c	22	18	1	13	54	197
Not satisfying extra constraints ^d	15	39	57	32	143	340
Poor light-curve fits ^e	2	24	8	28	62	402
Initial sample	172	335	236	279	1022	
Total discarded SNe	98	115	106	93	402	
Remaining SNe	74	230	130	186	620	

Notes. The reasons for discarding SNe are: (a) poor wavelength coverage that does not allow the calculation of the rest-frame B -band light curve, (b) poor time coverage that does not allow an initial estimation of t^{\max} , (c) unable to estimate an accurate B -band light curve after light-curve fit and correction, (d) unable to satisfy extra constraints on peak coverage (see the penultimate paragraph of Sec. 3), and (e) poor PISCOLA fits checked by visual inspection of SNe with large discrepancy in m_B^{\max} compared to SALT2 values.

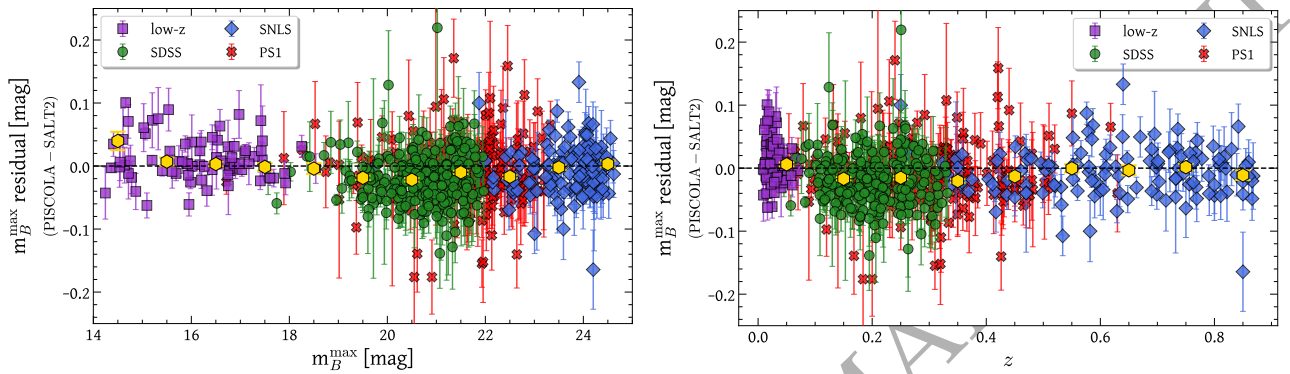


Figure 7. Comparison between m_B^{\max} obtained from the SALT2 light-curve fitter and that obtained using PISCOLA as a function of m_B^{\max} (left panel) and z (right panel) for SNe Ia from the Pantheon sample. The error bars are 1σ uncertainties from PISCOLA. The yellow hexagons represent the weighted mean in bins of 1 mag (left panel) and 0.1 (right panel) with their respective uncertainties (1σ).




Figure 8. Colour comparison between PISCOLA $(B - V)_{\max}$ and SALT2 c . The error bars are 1σ uncertainties from PISCOLA. Note that c is a more indirect measurement of $(B - V)_{\max}$ than the PISCOLA estimator.

5 SN IA LIGHT CURVE ANALYSIS

Having introduced PISCOLA as a general light-curve fitter tool for SNe, we now demonstrate a simple use case applied to the standardisation of SNe Ia for distance estimation. Many current and previous cosmological analyses (e.g., Astier et al. 2006; Guy et al. 2010; Betoule et al. 2014; Scolnic et al. 2018b) have used SALT/SALT2 light-curve

parametrizations to standardise SNe Ia using the Tripp-like formula (Tripp 1998):

$$\mu = m_B - M + \alpha \times x_1 - \beta \times c. \quad (2)$$

where μ is the SN Ia distance modulus, m_B , x_1 and c are the SALT2 light-curve parameters, and α , β and M are nuisance parameters. M represents the average B -band peak absolute magnitude of SNe Ia.

Current cosmological analyses using equation 2 incorporate an additional intrinsic dispersion term (σ_{int}), which encapsulates additional SN Ia variability that cannot be explained by their standardisation (e.g., Perlmutter et al. 1997; Tonry et al. 2003; Riess et al. 2004; Guy et al. 2010; Betoule et al. 2014; Scolnic et al. 2018b). In the following subsections we analyse SN Ia light curves from the Pantheon sample using Non-negative Matrix Factorization (NMF), an unsupervised machine-learning method, in search of an alternative parametrizations to study the standardisation of SNe Ia.

5.1 Light-Curve Decomposition with Non-negative Matrix Factorization

Previous works have used PCA (or variations of it) on SNe Ia to analyse their light curves and spectra to better understand their use as cosmological probes (e.g., Cormier & Davis 2011; Kim et al. 2013; Sasdelli et al. 2016;

He et al. 2018; Saunders et al. 2018). Here we use a different approach and analyse our rest-frame B -band SN Ia light curves with NMF, a data-driven linear decomposition algorithm, comparable to PCA. NMF factorises a non-negative matrix, \mathbf{X} , with dimensions $m \times n$, in a set of two non-negative matrices, \mathbf{W} and \mathbf{H} , with dimensions $m \times k$ and $k \times n$ (where $k < \min(n, m)$), respectively, such that $\mathbf{X} = \mathbf{W} \times \mathbf{H}$. In other words, NMF factorises a matrix into several components (eigen-vectors) and coefficients (eigen-values). Both methods (NMF and PCA) are commonly used for dimensionality reduction and feature extraction; however, NMF is better suited for some applications in astronomy as most astrophysical signals are non-negative, thus extracting components that are easier to interpret physically. Additionally, NMF components are not necessarily orthogonal, as is the case for PCA components, which allows us to find correlations between coefficients.

The light-curve decomposition depends on the phase coverage being used. We present our analysis for 214 SNe which have rest-frame data over the phase range -10 to $+15$ d. In this proof-of-concept investigation, we also discarded 50 SNe which did not have straight-forwardly rising and then declining light curves, flagged by visual inspection of the fits. There are many reasons why such fits may result, including astrophysical reasons such as the presence of secondary peaks in SN Ia light curves, or experimental reasons such as lower S/N data. However, this decomposition application is deliberately designed to be simple in scope; future work will examine the more complicated morphology of SN Ia light curves. The sample selection for this analysis is summarised in Table 5. In Section 5.3, we explain in more detail the phase range chosen, and explore other phase ranges. The rest-frame B -band light curves for this sample are shown in Fig. 9.

We decomposed these light curves in absolute-magnitude space, multiplied by -1 to obtain positive values. To calculate the absolute magnitudes we use distance moduli calculated with our assumed cosmology and the redshifts from Scolnic et al. (2018b, specifically z_{HD}). As the NMF algorithm used does not incorporate uncertainties, we used a Monte Carlo approach, generating sets of light curves for each SN from the uncertainties estimated by PISCOLA, and applying NMF decomposition to each of these sets. This generates a distribution of coefficients from which we used the mean value and the standard deviation to propagate uncertainties.

The following analysis is described for a decomposition with three components, although we consider further components in Section 5.3. In Fig. 10, we show the components obtained with NMF, with their respective explained variance (in percentage) with respect to the total variance of all components. Each eigen-vector contains specific information about the B -band of an ‘average’ SN Ia. Component 0 contributes to the general scale of the light curve, thus correlating with the B -band peak absolute magnitude, M_B^{max} . Component 1 contributes to the rise of the light curve, while component 2 mainly contributes to the decline of the light curve (and therefore we expect it to correlate with $\Delta m_{15}(B)$). Component 1 also contributes to the decline, but to a lesser degree than component 2. We note that from this data-driven decomposition, we naturally retrieve components related to the rise and decline of the B -band light

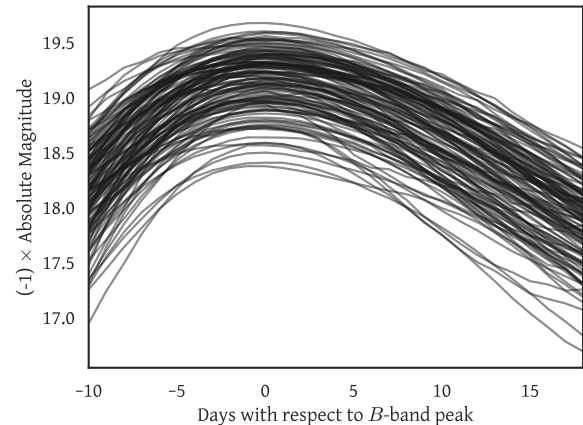


Figure 9. B -band light curves of 214 PISCOLA-fitted SNe Ia used in the decomposition analysis. Values are in absolute-magnitude multiplied by -1 .

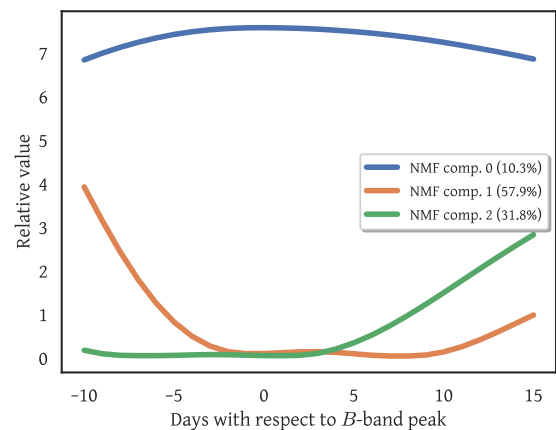


Figure 10. Three NMF components obtained from the decomposition of the B -band light curves shown in Fig. 9. In parenthesis are the percentages of the explained variance for each component with respect to the total variance of all components. Note that the absolute values of the components are not important for this analysis.

curve, in agreement with the findings of Hayden et al. (2010) and Ganeshalingam et al. (2011).

We are principally interested in the relative values of the components rather than the absolute values, which have no direct physical interpretation. We label the coefficients p_0 , p_1 and p_2 , associated to components 0, 1 and 2, respectively.

In Fig. 11, we show an example of a B -band light curve from one of the SNe together with a reconstructed light curve using the NMF components and coefficients, and their residuals. The residuals show that the reconstructed light curve has small differences ($\lesssim 0.05$ mag) to the original light-curve during the rise, but better agreement around peak (~ 0.02 mag residuals) and at later times (~ 0.00 mag residuals). The reconstructed light-curves and the original light-curves for the sample of 214 SNe are generally in excellent agreement, with mean residuals of ~ 0.0 mag and a standard

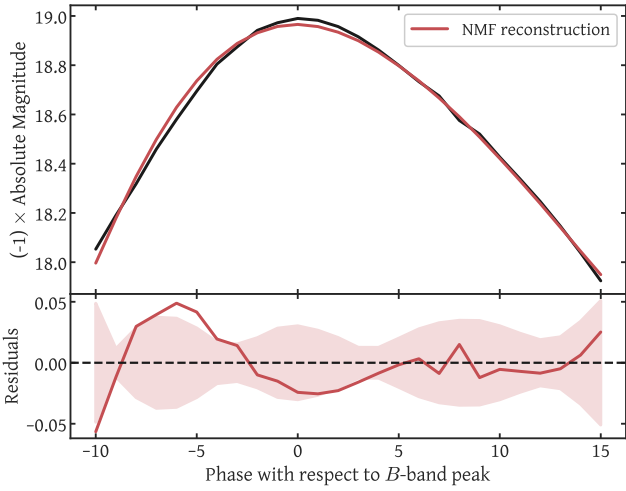


Figure 11. **Top:** example of B -band light curve from one of the SNe Ia (black) together with a reconstructed light-curve using the NMF components and coefficients (red). Values are in absolute-magnitude multiplied by -1 . **Bottom:** residuals from the light-curves shown in the top panel. The shaded area shows the standard deviation of the residuals for the entire sample in Fig. 9.

deviation of ~ 0.03 mag for all phases. This demonstrates that the NMF decomposition with three components is able to capture the variation in the light curves of SNe Ia.

The coefficients tell us about the contribution (or weight) of each of the components on the light curve of a SN Ia. By comparing these with different light-curve parameters, we can better understand their physical interpretation. This is shown in Fig. 12. As expected, there is a clear correlation between p_0 and M_B^{\max} , and thus $(B-V)_{\max}$ (e.g., Tripp 1998). Coefficient p_1 , which contributes to the rise of the light curve, has a small correlation with $(B-V)_{\max}$, but not with M_B^{\max} . On the other hand, coefficient p_2 clearly correlates with $\Delta m_{15}(B)$, as expected, and shows minor correlations with M_B^{\max} and $(B-V)_{\max}$. None of the components correlates with host-galaxy stellar mass. The decomposition is somewhat analogous to the SALT2 model, which contains x_0 and x_1 terms, although here we have two stretch components. In our case, the average (B -band light-curve) model has stretch parameters $p_1 \sim p_2 \sim 0.2$, while in the case of SALT2 it has $x_1 = 0$.

We note that NMF, unlike PCA, does not produce orthogonal components given the constraints of non-negative values, i.e., correlations between the components can be expected. From the coefficients, we find that p_1 and p_2 are anti-correlated (correlation of ~ -0.5).

5.2 Distance estimation

We follow an analogous approach to equation 2 using the PISCOLA components:

$$\mu = m_B - M + \eta_1 \times \hat{p}_1 + \eta_2 \times \hat{p}_2 - \beta \times (B - V)_{\max} \quad (3)$$

where m_B and $(B - V)_{\max}$ are the PISCOLA-measured light-curve parameters (i.e., measured from the fits), $\hat{p} =$




Figure 12. NMF coefficients vs PISCOLA light-curve parameters and host galaxy stellar mass for the SNe in Fig. 9. From the 214 SNe, only 156 have $(B - V)_{\max}$ values. Uncertainties are not shown for visualisation. Linear regressions are represented by red lines.

$p - \langle p \rangle$, and $\langle p \rangle$ is the average NMF coefficient value of the sample.




5.2.1 Hubble Diagram

To build a Hubble diagram, we adopt our standard cosmology (see Section 1) and use a Markov Chain Monte Carlo (MCMC) to find the optimal values for the nuisance parameters M , η_1 , η_2 and β , adopting uniform priors without bounds (except for M , where we require $M < 0$). Note that we used only SNe with $(B - V)_{\max}$ values (156 out of 214) for the Hubble diagram. See Table 5 for more details regarding the cuts applied to the SNe sample and the end of Section 5.3 for a summary. The need of colour limits the range of redshift, as we require rest-frame V -band coverage to estimate $(B - V)_{\max}$. Future work will focus on alternative colour parameters. For instance, the slope of the mangling function contains information about $(B - V)_{\max}$, which is approximately contained in a narrow range (~ 4500 – 5500 Å), allowing us to incorporate higher-redshift SNe, although it provides a more limited standardisation compared to $(B - V)_{\max}$.

We included uncertainties and covariances for our light-curve parameters, uncertainties in redshift (σ_z), uncertainties due to peculiar velocities (σ_{pec} ; adopting 300 km s^{-1}) and uncertainties from stochastic gravitational lensing (σ_{lens}). We also included an intrinsic scatter term (σ_{int}) such that $\chi_{\text{red}}^2 = 1$. Bias corrections are not included as a separate analysis would be required. However, we do not expect it to be significant due to the need for a $(B - V)_{\max}$ measurement in our analysis, which translates into an implicit redshift cut; see the redshift ranges for the different surveys in Fig. 13. The resulting Hubble diagram is shown in Fig. 13, while the MCMC results of the parameters are in Fig. 14.

Table 5. Number of SNe Ia discarded at different stages of the analysis (Section 5).

Discarding reason	Low-z	SDSS	SNLS	PS1	Total	Cumulative number discarded
Incomplete phase coverage	59	129	71	97	356	356
Unphysical light curve	1	19	8	22	50	406
No $(B - V)_{\max}$ estimation	0	12	40	6	58	464
Initial sample	74	230	130	186	620	
Total discarded	60	160	119	125	464	
Remaining SNe	14	70	11	61	156	


Figure 13. Hubble diagram (top panel) and residual (bottom panel) for the SNe Ia in the sample, using the PISCOLA standardisation introduced in this work.

Figure 14. MCMC results of the nuisance parameters used for the standardisation of the SNe Ia in the Hubble diagram (Fig. 13).

Figure 15. Hubble Residual comparison between the parametrization derived in this work (r.m.s. = 0.118 mag) and the standard parametrization from SALT2 (r.m.s. = 0.111 mag).

5.2.2 Hubble Residuals: PISCOLA vs SALT2

We now compare this simple PISCOLA parametrization with SALT2 using the same sample of 156 SNe. In Fig. 15, we show the comparison of the Hubble residuals measured using parameters derived from the PISCOLA light curves and from SALT2. PISCOLA obtains a similar r.m.s. value of 0.118 mag compared to SALT2 (0.111 mag), indicating it as a competitive method. The performance difference in r.m.s. of 0.007 mag is small. We note, however, that, although PISCOLA does not outperform SALT2, using it for cosmological analyses is a limited demonstration of its potential as a general purpose light-curve fitter.

The nuisance parameters obtained by using PISCOLA and SALT2 are summarised in Table 6. For SALT2, we found a similar value for α and a slightly lower value for β than those reported in Scolnic et al. (2018b), possibly due to the subsample used in this work. The value for M is different between both approaches, but as a normalisation factor this has no effect on the analysis. The value of β is smaller with PISCOLA, indicating the colour correction is not as large as with SALT2, possibly implying that some of the colour contribution is accounted for in p_1 and/or p_2 (note also the slight correlation between β and η_2 in Fig. 14). The stretch-like parameters, p_1 and p_2 , distribute as quasi-Gaussians over $\sim 0.0 - 0.4$ with mean values of ~ 0.2 , so the correction from $\eta_1 p_1 + \eta_2 p_2$ is larger than αx_1 . This is consistent with the possibility that some of the contribution from colour is absorbed by the p_1 and/or p_2 parameters. Finally, we obtain a lower σ_{int} with PISCOLA than with SALT2, but this is also a reflection of our larger uncertainties, where the main contributor to the uncertainty budget is $(B - V)_{\max}$, an effect of the GP light-curve fits.

Table 6. Nuisance parameters from the cosmological analysis for SALT2 and PISCOLA.

Parameter	SALT2	PISCOLA
M	$-19.36^{+0.01}_{-0.01}$	$-19.16^{+0.01}_{-0.01}$
β	$2.85^{+0.08}_{-0.08}$	$2.15^{+0.11}_{-0.12}$
α	$0.14^{+0.01}_{-0.01}$	-
η_1	-	$1.04^{+0.16}_{-0.15}$
η_2	-	$0.98^{+0.14}_{-0.13}$
σ_{int}	0.068	0.047

Notes. The uncertainties in σ_{int} are negligible (<0.001) in both cases.




Figure 16. PISCOLA Hubble residuals as a function of host galaxy stellar mass. Uncertainties are not shown for visualisation. A host galaxy mass step of -0.052 ± 0.022 mag (2.4σ) was obtained. The vertical dashed line marks the location of the step at $M_{\text{stellar}} = 10^{10} M_{\odot}$.

We also examined the dependence of Hubble residual on host galaxy stellar mass, the so-called ‘mass step’ (e.g., Kelly et al. 2010; Lampeitl et al. 2010; Sullivan et al. 2010), by using the stellar mass values from Scolnic et al. (2018b). We found a mass-step value of -0.052 ± 0.022 mag (2.4σ ; see Fig. 16) and -0.071 ± 0.016 mag (4.4σ) for PISCOLA and SALT2, respectively, for a step at $M_{\text{stellar}} = 10^{10} M_{\odot}$. The PISCOLA value is therefore consistent with the value obtained with SALT2 and with previous work (e.g., Sullivan et al. 2010; Betoule et al. 2014; Scolnic et al. 2018b; Kelsey et al. 2021; Boone et al. 2021), although lower and less significant.

5.3 Further Exploration

The results of the light-curve decomposition in Section 5.1 depend on the phase range considered. We explored different ranges with lower limits of -8 , -10 and -12 days, and upper limits of $+12$, $+15$ and $+18$ days. As not all SNe have the same coverage, we used an initial sample of 264 SNe in common for all the ranges. Data outside of these phase ranges are usually more incomplete. 50 SNe were discarded for having unphysical-looking light curves, as previously described, leaving 214 SNe. From these 214 SNe, only 156 SNe had $(B - V)_{\text{max}}$ values. For details about the discarded SNe at different stages of the analysis process, see Table 5.

Our full analysis (Sections 5.1 and 5.2) was performed for these phase ranges and the Hubble residual r.m.s. values

Table 7. Hubble residual (HR) r.m.s. for our method using different combinations of light-curve phase ranges. The range that gave the lowest r.m.s. is in bold.

Minimum phase [days]	Maximum phase [days]	HR r.m.s. [mag]
-8	+12	0.125
-8	+15	0.120
-8	+18	0.123
-10	+12	0.123
-10	+15	0.118
-10	+18	0.122
-12	+12	0.125
-12	+15	0.125
-12	+18	0.127

Notes. The HR r.m.s. for the same sample, using SALT2, is 0.111 mag.

are in Table 7. The range of $[-10, +15]$ d produced the smallest r.m.s. Although naively it might be expected that larger ranges would contain more light-curve information in the NMF components, the fits can also be less reliable at early or late epochs. For example, the ranges starting at -12 days produce the largest r.m.s. values, due to less reliable fits to the larger observational uncertainties at these early epochs. A similar behaviour is seen for $+18$ d. The range of $[-10, +15]$ d produces the best combination between information incorporated and reliable fits.

The shape and information contained in the different NMF components also depends on the number of components chosen. We repeated the analysis using two, four and five components for the phase range of $[-10, +15]$ d. Using two components resulted in a clearly worse result (Hubble residual r.m.s. of 0.141 mag), while using four and five components produced similar results as that with three components (both with a Hubble residual r.m.s. of 0.121 mag). We conclude that two components are not sufficient to capture the variability of SNe Ia. On the other hand, as the number of components increases it is hard to determine if there is any physical interpretation and/or contribution to the light-curve standardisation. The results suggest that three components is optimal with the current data, each component with clear physical interpretation. Although we did not produce an improved performance over SALT2, despite using one extra parameter, we stress that this is an exploratory work where a single-band light-curve decomposition was performed. Future work with multiple bands may produce improved standardisation as stretch and colour information would be included.

To summarise the sample used for the cosmological analysis, of the 620 SNe that PISCOLA successfully fit (obtained m_B^{max}), only 264 have phase coverage in the B -band between -12 and $+18$ days. There are more SNe with smaller phase-range coverage, but to ensure a fair comparison, we use the sample from the largest phase-range coverage we tested, which has SNe in common with all the other phase ranges. Of these 264 SNe, 50 had unphysical-looking light curves from visual inspection and were removed, leaving 214 SNe on which the NMF analysis was formed. Finally, from these 214 SNe, only 156 have $(B - V)_{\text{max}}$ and form our final cosmological sample. We note that, although the cuts in the sample are severe, this is a proof-of-concept analysis with

the idea of exploring new approaches and thus requires a ‘golden’ sample, removing possible biases.

6 PISCOLA COLOUR LAW

The SN colour is the light-curve parameter that contributes most to the standardisation of SNe Ia. It is related to the physics, progenitors and environments of SNe Ia, and is a complicated parameter as there are several factors that may contribute to variations in the observed colour (e.g., SN circumstellar material or host galaxy extinction).

The wavelength-dependent variation of colour in SNe Ia is known as the ‘colour-variation law’ (CL). SALT2 describes the CL as a wavelength-dependent function that does not vary in time or as a function of x_1 (see equation (1)). We can also use the PISCOLA mangling fits to estimate a CL, in a similar fashion to SALT2. The mangling function at t^{\max} (see Section 2.2) describes how the colour of a SN varies with respect to a base SED and therefore gives us a method to estimate the relative CL for a single SN (i.e., the mangling function is equal to $(B - V)_{\max}$ multiplied by the CL).

We divide SNe into bins of 0.05 mag in colour, and calculate their average mangling function in each bin. We then fit a third-order polynomial, optimising across all bins of $(B - V)_{\max}$ simultaneously, to obtain a functional form for the CL following the assumption that the CL is wavelength-dependent only (as in equation (1)). Only data between ~ 3500 and ~ 7000 Å is used as not all SNe have coverage bluer than 3500 Å, especially at low redshift. High- z SNe get their ultraviolet (UV) wavelengths redshifted to optical wavelengths in the observer frame. However, most of the SNe with rest-frame UV coverage do not have PISCOLA-measured $(B - V)_{\max}$ as the rest-frame V band gets redshifted outside the filters’ range, and so are not used for the estimation of the PISCOLA CL.

In Fig. 17, we compare the PISCOLA CL against that of SALT2 and the extinction law from Fitzpatrick (1999) for three different values of the total to selective extinction ratio, R_V . The PISCOLA CL agrees with the Fitzpatrick (1999) extinction law with $R_V \lesssim 3.1$. Similar findings have been reported in previous work (e.g., Burns et al. 2014; Amanullah et al. 2015; Sasdelli et al. 2016; Brout & Scolnic 2021; Thorp et al. 2021).

When comparing with the SALT2 CL, there is agreement at optical wavelengths, but some deviation towards the UV. However, as previously mentioned, the data around 3000 Å is limited and thus the PISCOLA CL at these wavelengths is more uncertain.

The colour dispersion, i.e., the scatter around the CL, is also of importance to quantify the disagreement. SALT2 includes its CL in its model (equation (1)), possibly limiting its behaviour, and it is estimated during the training phase, while we estimate it after the PISCOLA fitting and correction process of the SNe Ia, almost directly from the data. In Fig. 18, we compare the CLs of PISCOLA and SALT2 at $(B - V)_{\max} = 0.15$ mag, but also include their respective colour dispersions. We note that, despite the differences in the estimation of the CLs, they agree within the uncertainties, especially in the range used for calculating the PISCOLA CL (at ~ 3500 – 7000 Å).

This was also tested with the simulations of Sec. 3.1,

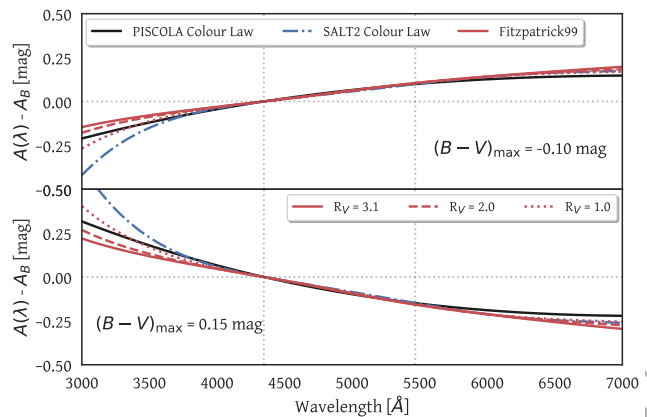


Figure 17. The colour law obtained with PISCOLA (solid black line) compared to that from SALT2 (dash-dotted blue line) and a Fitzpatrick (1999, Fitzpatrick99) extinction law for different R_V values (red lines). The vertical dotted lines mark the effective wavelengths of the B and V bands. This comparison is shown for two different $(B - V)_{\max}$ values (-0.10 and 0.15 mag).

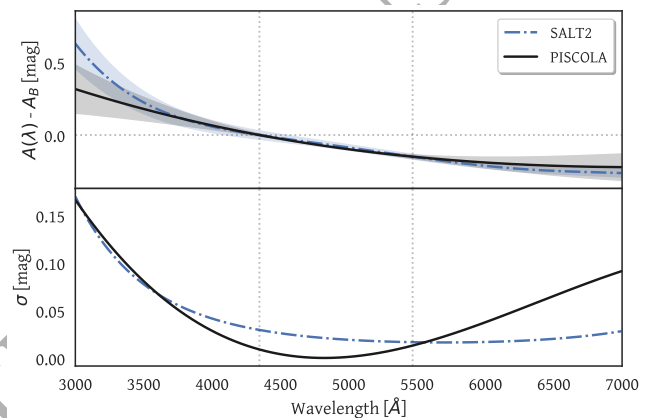


Figure 18. **Top panel:** CL obtained with PISCOLA (solid black line) compared to the CL from SALT2 (dash-dotted blue line) for $(B - V)_{\max} = 0.15$ mag. The shaded areas represent their respective colour dispersions. **Bottom panel:** Colour dispersion comparison between PISCOLA (solid black line) and SALT2 (dash-dotted blue line). The vertical dotted lines mark the effective wavelengths of the B and V bands.

where we found similar results. Furthermore, by changing the shape of the CL of the simulations, PISCOLA is able to retrieve a CL that agrees with it (at ~ 3500 – 7000 Å; see Fig. 19).

The colour dispersion indicates that PISCOLA has a larger scatter towards the near-IR, but similar in the UV, compared to SALT2. The shape of the colour dispersion curve from PISCOLA could be a consequence of a combination of three factors. The first is due to the smaller observational uncertainties in the B and V bands, where much of the flux of a SN Ia emerges, and larger uncertainties at redder and bluer wavelengths. The second is due to the reduced number of SNe covering bluer and redder wavelengths (around 3500 and 7000 Å respectively) compared to the range around B and V . The third is that there may be a difference in the CL for different SNe.

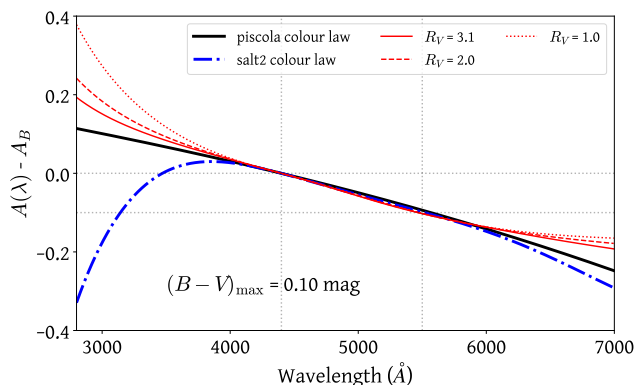


Figure 19. The description is the same as in Fig. 17 but for $(B - V)_{\max} = 0.10$ mag and a SALT2 model with a different (artificial) CL shape. PISCOLA is able to retrieve a CL that agrees with that of SALT2 (at $\gtrsim 3500$ Å) even after changing the shape of the CL in the SALT2 model.

Differences in CLs are important as they can represent differences in the underlying physics of SNe Ia. The SALT2 model mainly incorporates near-UV observations of high- z SNe (e.g., from SNLS) as these wavelengths are redshifted to optical wavelengths and can be observed with ground-based telescopes. In principle, the characteristics of these SNe could be intrinsically different to those at lower redshift (e.g., Ellis et al. 2008; Maguire et al. 2012). This may introduce biases in the SALT2-measured colour. The CL obtained with PISCOLA agrees with the average extinction law measured in the MW (with $R_V \lesssim 3.1$) within the uncertainties, implying that the remaining variation (intrinsic scatter) in SNe Ia may be driven by dust of similar properties to that in our galaxy. We note that Brout & Scolnic (2021) found the opposite, i.e., dust with different properties to that in our galaxy, although using a SALT2 CL.

7 SUMMARY

In this work we have presented a new open-source SN light-curve fitting code PISCOLA, which relies on Gaussian processes (GP), a data-driven interpolation method, for fitting light-curves. PISCOLA can be applied to any observer-frame SN light-curves, and produces rest-frame light-curves as its principal output. PISCOLA can additionally estimate rest-frame light-curve parameters, such as peak magnitudes, colours, and light-curve shapes.

We tested PISCOLA by applying it to SN Ia data, both simulated and real. With simulations of SNe Ia for different cadences and observational uncertainties, we found that PISCOLA is reliable for observational cadences of $\lesssim 7$ days for typical current SN Ia samples, provided relatively loose constraints on data coverage around peak luminosity and signal-to-noise are used. When comparing PISCOLA outputs on real data and comparing to light-curve fits with the SALT2 light-curve fitter, we see small but significant ($>3\sigma$) differences in peak rest-frame B -band magnitude. However, with no ground-truth for these tests, we argue that such differences may be expected given the different assumptions used in the two techniques.

We then demonstrated a scientific use of PISCOLA by

analysing the rest-frame B -band light curves of the Pantheon SN Ia sample using NMF, a machine-learning decomposition algorithm, to search for alternative standardisations of these objects. NMF allows the extraction of ‘easy-to-interpret’ and non-orthogonal components, unlike other algorithms, such as PCA. We compared the NMF coefficients with different SN Ia parameters, and used them to build a Hubble diagram. We tested different combinations of light-curve phase ranges and numbers of components for the decomposition, and found the best results were based on B -band light curves with a phase range of $[-10, +15]$ d and three components. This parametrization produces an r.m.s. in the Hubble residual similar to that of SALT2 (0.118 and 0.111 mag, respectively), showing the promise of this new framework. We stress that, although PISCOLA does not outperform SALT2 in this particular case, future work with multiple bands can produce better results. Additionally, this analysis is only a limited demonstration of PISCOLA’s potential.

PISCOLA uses a smooth GP interpolation to adjust its base SED to an observed SN colour (a mangling function). This mangling function encodes information on the colour law (CL) of SNe Ia. We estimated a functional form for this CL by fitting a third-order polynomial and compared it with the SALT2 CL and Fitzpatrick (1999) extinction laws with different R_V values. We found that the PISCOLA CL agrees with an extinction law with $R_V \lesssim 3.1$, but also with the SALT2 CL. Although there could be some slight disagreement towards the UV, a possible cause is the extrapolation of the PISCOLA CL bluer than ~ 3500 Å.

We have plans for future upgrades of PISCOLA. This includes the use of a ‘stretch’-dependent time-series SED, a mangling function in two-dimensions (wavelength and time), and exploring alternative GP models, for example the use of different kernels, different bounds for the hyper-parameters, etc.. This may produce more accurate light-curve fits and mangling functions. Finally, we emphasise that although our tests have been based around applications to SNe Ia, PISCOLA is generic so that, with an appropriate time-series SED for K -corrections, it can fit any type of optical transient to estimate rest-frame light-curves and luminosities. For instance, SED templates for type Ib, Ic and II SNe already exist (e.g., Vincenzi et al. 2019) and, in the case of superluminous SNe, a black-body could be used as an approximate SED template as no current templates exist.

ACKNOWLEDGEMENTS

Software:

numpy (Harris et al. 2020), matplotlib (Hunter 2007), seaborn (Waskom et al. 2017), pandas (McKinney et al. 2010), scipy (Virtanen et al. 2020), lmfit (Newville et al. 2016), scikit-learn (Pedregosa et al. 2011), emcee (Foreman-Mackey et al. 2013), coner (Foreman-Mackey 2016), george (Ambikasaran et al. 2016), sfdmap⁴, extinction⁵ (Barbary 2016), peakutils⁶ (Negri & Vestri 2017), snocosmo (Barbary et al. 2016).

⁴ <https://github.com/kbarbary/sfdmap>

⁵ <https://github.com/kbarbary/extinction>

⁶ <https://github.com/lucashn/peakutils/tree/v1.1.0>

TMB acknowledges funding from the CONICYT PFCCHA/DOCTORADOBECCAS CHILE/2017-72180113. We thank Dalya Baron for discussions regarding the use and understanding of principal component analysis and similar algorithms. TMB would also like to thank the LSSTC-DSFP which allowed him to develop computational skills and tools for the development of PISCOLA and the analysis in this work. We thank Santiago González-Gaitán and Lluís Galbany for discussions about light-curve fitting and decomposition. Finally, we thank Dan Scolnic for his patience in helping us understand the intricacies of the calibration of the Pantheon SN Ia sample.

DATA AVAILABILITY

All the data used in this work was publicly available. PISCOLA is an open source code and can be found in a Github repository: <https://github.com/temuller/piscola>.

REFERENCES

- Abbott T. M. C., et al., 2019, *ApJ*, 872, L30
- Amanullah R., et al., 2015, *MNRAS*, 453, 3300
- Ambikasaran S., Foreman-Mackey D., Greengard L., Hogg D. W., O’Neil M., 2016, *IEEE Transactions on Pattern Analysis and Machine Intelligence*, 38, 252
- Angus C. R., et al., 2019, *MNRAS*, 487, 2215
- Astier P., et al., 2006, *A&A*, 447, 31
- Barbary K., 2016, extinction v0.3.0, doi:10.5281/zenodo.804967, <https://doi.org/10.5281/zenodo.804967>
- Barbary K., et al., 2016, sncosmo/sncosmo: v1.4.0, doi:10.5281/zenodo.168220, <https://doi.org/10.5281/zenodo.168220>
- Bessell M. S., 1990, *PASP*, 102, 1181
- Betoule M., et al., 2014, *A&A*, 568, A22
- Blondin S., Tonry J. L., 2007, *ApJ*, 666, 1024
- Bohlin R. C., Gilliland R. L., 2004, *AJ*, 128, 3053
- Boone K., et al., 2021, arXiv e-prints, p. arXiv:2105.02204
- Brout D., Scolnic D., 2021, *ApJ*, 909, 26
- Burns C. R., et al., 2011, *The Astronomical Journal*, 141, 19
- Burns C. R., et al., 2014, *The Astrophysical Journal*, 789, 32
- Cadonau R., Sandage A., Tammann G. A., 1985, Type I Supernovae as Standard Candles. p. 151, doi:10.1007/3-540-15206-7_56
- Colgate S. A., McKee C., 1969, *The Astrophysical Journal*, 157, 623
- Conley A., et al., 2008, *The Astrophysical Journal*, 681, 482
- Conley A., et al., 2011, *The Astrophysical Journal Supplement Series*, 192, 1
- Contreras C., et al., 2010, *The Astronomical Journal*, 139, 519
- Cormier D., Davis T. M., 2011, *MNRAS*, 410, 2137
- Ellis R. S., et al., 2008, *ApJ*, 674, 51
- Filippenko A. V., et al., 1992a, *AJ*, 104, 1543
- Filippenko A. V., et al., 1992b, *ApJ*, 384, L15
- Fitzpatrick E. L., 1999, *PASP*, 111, 63
- Folatelli G., et al., 2010, *AJ*, 139, 120
- Foreman-Mackey D., 2016, *The Journal of Open Source Software*, 1, 24
- Foreman-Mackey D., Hogg D. W., Lang D., Goodman J., 2013, *PASP*, 125, 306
- Frieman J. A., et al., 2008, *AJ*, 135, 338
- Ganeshalingam M., Li W., Filippenko A. V., 2011, *MNRAS*, 416, 2607
- Graur O., et al., 2014, *ApJ*, 783, 28
- Guy J., et al., 2007, *Astronomy and Astrophysics*, 466, 11
- Guy J., et al., 2010, *Astronomy and Astrophysics*, 523, A7
- Harris C. R., et al., 2020, *Nature*, 585, 357
- Hayden B. T., et al., 2010, *ApJ*, 712, 350
- He S., Wang L., Huang J. Z., 2018, *ApJ*, 857, 110
- Hicken M., et al., 2009a, *The Astrophysical Journal*, 700, 331
- Hicken M., Wood-Vasey W. M., Blondin S., Challis P., Jha S., Kelly P. L., Rest A., Kirshner R. P., 2009b, *ApJ*, 700, 1097
- Hicken M., et al., 2012, *ApJS*, 200, 12
- Hoyle F., Fowler W. A., 1960, *ApJ*, 132, 565
- Hsiao E. Y., Conley A., Howell D. A., Sullivan M., Pritchett C. J., Carlberg R. G., Nugent P. E., Phillips M. M., 2007, *ApJ*, 663, 1187
- Hunter J. D., 2007, *Computing in Science Engineering*, 9, 90
- Inserra C., et al., 2018, *MNRAS*, 475, 1046
- Jha S., et al., 2006, *The Astronomical Journal*, 131, 527
- Jha S., Riess A. G., Kirshner R. P., 2007, *ApJ*, 659, 122
- Jha S. W., Maguire K., Sullivan M., 2019, *Nature Astronomy*, 3, 706
- Kattner S., et al., 2012, *Publications of the Astronomical Society of the Pacific*, 124, 114
- Kelly P. L., Hicken M., Burke D. L., Mandel K. S., Kirshner R. P., 2010, *ApJ*, 715, 743
- Kelsey L., et al., 2021, *MNRAS*, 501, 4861
- Kessler R., et al., 2009a, *PASP*, 121, 1028
- Kessler R., et al., 2009b, *ApJS*, 185, 32
- Kessler R., et al., 2013, *ApJ*, 764, 48
- Kim A. G., et al., 2013, *ApJ*, 766, 84
- Kowal C. T., 1968, *AJ*, 73, 1021
- Lampeitl H., et al., 2010, *MNRAS*, 401, 2331
- Landolt A. U., 1992, *AJ*, 104, 340
- Leibundgut B., Tammann G. A., Cadonau R., Cerrito D., 1991, *A&AS*, 89, 537
- Liu Y., Modjaz M., 2014, arXiv e-prints, p. arXiv:1405.1437
- Lochner M., et al., 2018, arXiv e-prints, p. arXiv:1812.00515
- Maguire K., et al., 2012, *MNRAS*, 426, 2359
- McKinney W., et al., 2010, in *Proceedings of the 9th Python in Science Conference*. pp 51–56
- Negri L. H., Vestri C., 2017, lucashn/peakutils: v1.1.0, doi:10.5281/zenodo.887917, <https://doi.org/10.5281/zenodo.887917>
- Newville M., Stensitzki T., Allen D. B., Rawlik M., Ingargiola A., Nelson A., 2016, Lmfitt: Non-Linear Least-Square Minimization and Curve-Fitting for Python (ascl:1606.014)
- Nugent P., Kim A., Perlmutter S., 2002, *PASP*, 114, 803
- Pedregosa F., et al., 2011, *Journal of Machine Learning Research*, 12, 2825
- Perlmutter S., et al., 1997, *ApJ*, 483, 565
- Perlmutter S., et al., 1999, *ApJ*, 517, 565
- Phillips M. M., 1993, *ApJ*, 413, L105
- Phillips M. M., Lira P., Suntzeff N. B., Schommer R. A., Hamuy M., Maza J., 1999, *AJ*, 118, 1766
- Pskovskii I. P., 1977, *Soviet Ast.*, 21, 675
- Pursiainen M., et al., 2020, *MNRAS*, 494, 5576
- Rasmussen C. E., Williams C. K. I., 2006, *Gaussian Processes for Machine Learning*
- Rest A., et al., 2014, *ApJ*, 795, 44
- Riess A. G., Press W. H., Kirshner R. P., 1995, *ApJ*, 438, L17
- Riess A. G., Press W. H., Kirshner R. P., 1996, *ApJ*, 473, 88
- Riess A. G., et al., 1998, *AJ*, 116, 1009
- Riess A. G., et al., 1999a, *AJ*, 117, 707
- Riess A. G., et al., 1999b, *AJ*, 118, 2675
- Riess A. G., et al., 2004, *ApJ*, 607, 665
- Riess A. G., et al., 2007, *ApJ*, 659, 98
- Riess A. G., et al., 2011, *ApJ*, 730, 119
- Riess A. G., et al., 2018, *ApJ*, 853, 126
- Rodney S. A., Tonry J. L., 2009, *ApJ*, 707, 1064
- Rodney S. A., et al., 2014, *AJ*, 148, 13

- Ruiz-Lapuente P., et al., 1993, *Nature*, **365**, 728
 Sako M., et al., 2018, *PASP*, **130**, 064002
 Sasdelli M., Ishida E. E. O., Hillebrandt W., Ashall C., Mazzali P. A., Prentice S. J., 2016, *MNRAS*, **460**, 373
 Saunders C., et al., 2018, *ApJ*, **869**, 167
 Schlafly E. F., Finkbeiner D. P., 2011, *ApJ*, **737**, 103
 Schlegel D. J., Finkbeiner D. P., Davis M., 1998, *ApJ*, **500**, 525
 Scolnic D., Kessler R., 2016, *ApJ*, **822**, L35
 Scolnic D., et al., 2014, *ApJ*, **795**, 45
 Scolnic D. M., et al., 2018a, arXiv e-prints, p. [arXiv:1812.00516](https://arxiv.org/abs/1812.00516)
 Scolnic D. M., et al., 2018b, *ApJ*, **859**, 101
 Sollerman J., et al., 2009, *ApJ*, **703**, 1374
 Stritzinger M. D., et al., 2011, *AJ*, **142**, 156
 Sullivan M., et al., 2010, *MNRAS*, **406**, 782
 Suzuki N., et al., 2012, *ApJ*, **746**, 85
 Tammann G. A., Leibundgut B., 1990, *A&A*, **236**, 9
 Thorp S., Mandel K. S., Jones D. O., Ward S. M., Narayan G., 2021, arXiv e-prints, p. [arXiv:2102.05678](https://arxiv.org/abs/2102.05678)
 Tonry J. L., et al., 2003, *ApJ*, **594**, 1
 Tripp R., 1998, *A&A*, **331**, 815
 Vincenzi M., Sullivan M., Firth R. E., Gutiérrez C. P., Frohmaier C., Smith M., Angus C., Nichol R. C., 2019, *MNRAS*, **489**, 5802
 Virtanen P., et al., 2020, *Nature Methods*,
 Waskom M., et al., 2017, *mwaskom/seaborn*: v0.8.1 (September 2017), [doi:10.5281/zenodo.883859](https://doi.org/10.5281/zenodo.883859), <https://doi.org/10.5281/zenodo.883859>
 Wiseman P., et al., 2020, *MNRAS*, **498**, 2575
 Woosley S. E., Taam R. E., Weaver T. A., 1986, *ApJ*, **301**, 601
 de Jaeger T., et al., 2017, *MNRAS*, **472**, 4233

APPENDIX A: GAUSSIAN PROCESS

A GP is a random process where any point $\mathbf{x} \in \mathbb{R}^d$ is assigned a random variable $f(\mathbf{x})$ and where the joint distribution of a finite number of these variables $p(f(\mathbf{x}_1), \dots, f(\mathbf{x}_N))$ is itself Gaussian: $p(\mathbf{f} | \mathbf{X}) = \mathcal{N}(\mathbf{f} | \boldsymbol{\mu}, \mathbf{K})$, where $\boldsymbol{\mu}$ and \mathbf{K} are the mean and covariance functions. The latter is also called *kernel*.

PISCOLA incorporates three different kernels: Squared Exponential, Matérn-3/2 and Matérn-5/2. These are defined as

$$k_{SE}(x, x') = \sigma^2 \exp\left(-\frac{|x - x'|^2}{2\ell}\right), \quad (\text{A1})$$

$$k_{M32}(x, x') = \sigma^2 \left(1 + \frac{\sqrt{3}|x - x'|}{\ell}\right) \exp\left(-\frac{\sqrt{3}|x - x'|}{\ell}\right), \quad (\text{A2})$$

and

$$k_{M52}(x, x') = \sigma^2 \left(1 + \frac{\sqrt{5}|x - x'|}{\ell} + \frac{5|x - x'|^2}{3\ell^2}\right) \exp\left(-\frac{\sqrt{5}|x - x'|}{\ell}\right), \quad (\text{A3})$$

where ℓ (length-scale) and σ^2 (variance) are the hyperparameters of the kernels. Fig. A1 shows a comparison between the three kernels. From the mathematical perspective, each kernel has its own properties; however, in terms of applications for fits, the main difference is the covariance between points, which is larger for the Squared Exponential kernel than for the other two, assuming two data points at a fixed distance, and a fixed set of hyperparameters. The *Matérn*




Figure A1. Comparison of the covariance between three kernels available with PISCOLA: Squared Exponential (blue), Matérn-3/2 (orange) and Matérn-5/2 (green).

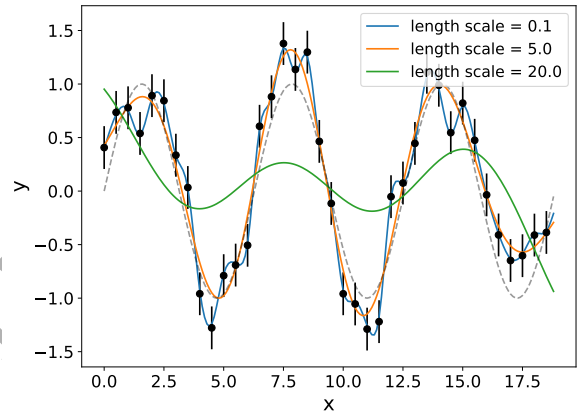


Figure A2. Comparison between different length scale values (ℓ) for an Squared Exponential kernel: 0.1 (blue), 5.0 (orange) and 20.0 (green). GP fits (solid lines) were performed on simulated data (black circles). Only the mean GP model is shown in each case for visualisation. The underlying function, from which the simulated data was extracted from, is also shown (dashed grey line).

family is appropriate for modelling physical processes (the Squared Exponential kernel can be too smooth), with the Matérn-3/2 and Matérn-5/2 kernels of special interest as they are simple functions.

Figs. A2 and A3 show a comparison between different length scales and variances for the Squared Exponential kernel. The length scale determines the length of the ‘wiggles’ in a function, giving a sense of how far the data can be extrapolated. The variance determines the average distance of the function away from its mean. For a detailed introduction to GP method see [Rasmussen & Williams \(2006\)](#).




Figure A3. Comparison between different variance values (σ^2) for an Squared Exponential kernel: 0.01 (blue), 0.1 (orange) and 10.0 (green). GP fits (solid lines) were performed on simulated data (black circles). Only the mean GP model is shown in each case for visualisation purposes. The underlying function, from which the simulated data was extracted from, is also shown (dashed grey line).

A1 One dimensional Gaussian process fits

Two-dimensional GP light-curve fits have the advantage of performing informative interpolation/extrapolation given data from multiple bands and this forms our default approach. However, we also explored fitting each filter independently (a 1D GP model). In Fig. A4, we show the same example as in Fig. 2, but for such a 1D GP model, with independent fits for each band, i.e. a regression of flux as a function of time. The extrapolation of the rise of the light curve for the r and z bands over-predicts the flux by comparing them with i band, given that the former bands are bluer and redder than the latter. In addition, the uncertainty of the z light-curve decline fit rapidly increases the further the GP model extrapolates. A final consideration is the estimation of t^{\max} : as in a 1D fit there is no interpolation in the wavelength axis, the peak can only be estimated by choosing the closest observer-frame band to rest-frame B band or using a redder and bluer band than the B band for a more accurate estimation. Both approaches are less precise than the estimation from a two-dimensional interpolation.

A2 Effect of S/N

The results of a GP fit depends on the cadence and uncertainties of the observations. In Fig. A5 we show a GP fit to a rest-frame B -band light curve taken from a SN Ia SED from which we extracted ‘observations’ with a 10-day cadence. We simulate fractional errors of the observations with 5, 10, 20 and 30 per cent values, reproducing the effect of S/N. When the errors are small (5 and 10 per cent), the peak is over-predicted (e.g., for low- z SNe). Conversely, when they are large (30 per cent), the peak is under-predicted (e.g., for SDSS SNe).

This just gives a general idea on how S/N can affect the light-curve fits as the exact values of the uncertainties will depend on the apparent magnitude of the SN. For example,




Figure A4. The description is the same as in Fig. 2, but for a 1D GP model, with independent fits for each band.




Figure A5. 10-day cadence ‘observations’ (red circles) extracted from the rest-frame B -band light curve (dashed grey line) of a SN Ia SED template. GP fits (in green) were performed for cases with simulated fractional errors of the observations of 5, 10, 20 and 30 per cent.

for two SNe with peak apparent magnitudes of 17 and 14, a 1 per cent error turns into 0.17 mag and 0.14 mag, respectively.

APPENDIX B: LIGHT-CURVE FITS

PISCOLA presents several advantages over template-driven light-curve fitters, one of them being the better fits of low- z SNe, i.e., objects with relatively well-sampled light curves and high S/N. One example is presented in Fig. B1, where we show the PISCOLA fit to SN 2004ey from the CSP survey. In Fig. B2, we show the SALT2 fit we performed to the same SN, using the implementation of the code in `sncosmo`. In this case, we see that SALT2 produces larger residuals compared to PISCOLA.

PISCOLA outperforms SALT2 when fitting the CSP u and r , which are on the blue and red limits, respectively, of the effective training range of SALT2 (~ 3000 – 7000 Å). Furthermore, PISCOLA produces good results in the i band, even fitting the secondary peak, showing the potential of

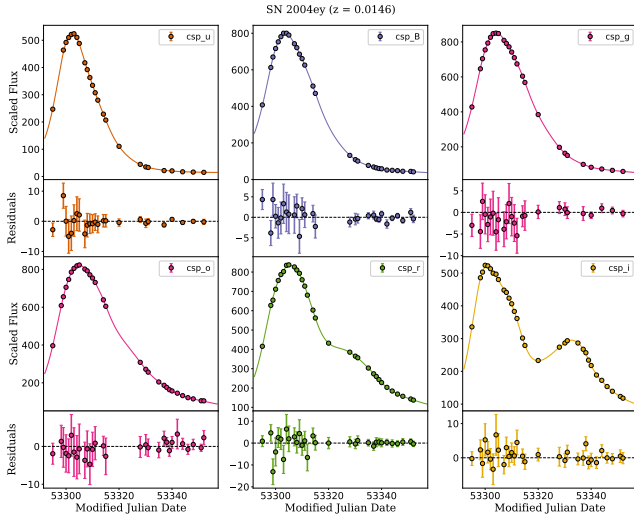


Figure B1. PISCOLA light-curves fit and residuals of SN 2004ey from the CSP survey. The observational uncertainties for this SN are small so they are only seen in the residual plots. The label `csp_o` is used by PISCOLA to refer to one of the three available CSP V-bands.

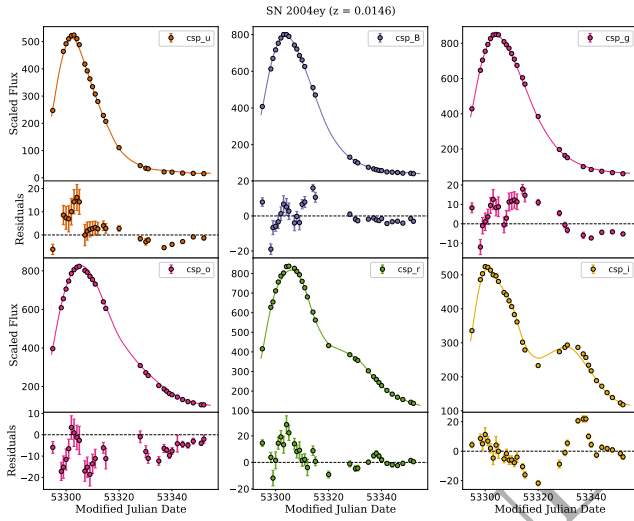


Figure B2. SALT2 light-curves fit and residuals of SN 2004ey from the CSP survey. The rest of the description is the same as in Fig. B1.

fitting NIR light curves. We also note that PISCOLA produces better results in bands like CSP *g* and *V* (`csp_o` in Fig. B2), important for the estimation of colour and colour evolution in SNe Ia.

Future work will involve using PISCOLA to fit NIR bands of SNe Ia and estimate distances with them.

This paper has been typeset from a TeX/L^AT_EX file prepared by the author.